

Spectral Ringing Artifacts in Hyperspectral Image Data Compression

M. Klimesh,¹ A. Kiely,¹ H. Xie,¹ and N. Aranki²

When a three-dimensional wavelet decomposition is used for compression of hyperspectral images, spectral ringing artifacts can arise, manifesting themselves as systematic biases in some reconstructed spectral bands. More generally, systematic differences in signal level in different spectral bands can hurt compression effectiveness of spatially low-pass subbands. The mechanism by which this occurs is described in the context of ICER-3D, a hyperspectral imagery extension of the ICER image compressor. Methods of mitigating or eliminating the detrimental effects of systematic band-dependent signal levels are proposed and discussed, and results are presented.

I. Introduction

Hyperspectral images are three-dimensional (3-D) data sets, where two of the dimensions are spatial and the third is spectral. A hyperspectral image can be regarded as a stack of individual images of the same spatial scene, with each such image representing the scene viewed in a narrow portion of the electromagnetic spectrum. These individual images are referred to as spectral bands. Hyperspectral images typically consist of more than 200 spectral bands; the voluminous amount of data comprising hyperspectral images makes them appealing candidates for data compression.

Straightforward extension of wavelet-based two-dimensional (2-D) image compression to hyperspectral image compression based on a 3-D wavelet decomposition can result in inefficient coding of some subbands and can lead to reconstructed spectral bands with systematic biases. In this article, we describe this problem in detail and discuss some methods to resolve it.

The effects we describe are consequences of the fact that the wavelet transform doesn't account for systematic differences in signal level in different spectral bands. We remark that using a wavelet transform for spectral decorrelation of hyperspectral data has other shortcomings as well. For example, the spectral dependencies that exist are not limited to the small spectral neighborhood exploited by the wavelet transform. However, the 3-D wavelet transform has practical advantages compared to other transforms: it offers reasonably effective compression with modest computational and implementation complexity.

¹ Communications Architectures and Research Section.

² Flight System Electronics Section.

The research described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

Our analysis and results are presented with respect to the ICER-3D compressor, which was created as an extension of the ICER image compressor to hyperspectral images. ICER is a wavelet-based, progressive (embedded), 2-D image compressor; see [1] for a description. ICER is being used onboard the Mars Exploration Rovers for compression of a large majority of the images returned [2]. ICER-3D inherits much of its design from ICER, but uses a 3-D wavelet decomposition to provide decorrelation in the spectral dimension as well as both spatial dimensions. Further development of ICER-3D is ongoing.

Other investigations of 3-D wavelet-based compression of hyperspectral imagery include [3–5].

The examples presented in this article use Airborne Visible/Infrared Imaging Spectrometer (AVIRIS) data [6]. AVIRIS hyperspectral images have a width of 614 pixels and include 224 spectral bands covering wavelengths from 370 nm to 2500 nm. For most of our examples, we use the first 512 line scene of the calibrated 1997 Moffett Field radiance data set.³ In this article, we number bands, columns, rows, and planes starting from 1 (rather than 0).

For all examples in this article, wavelet transforms are performed using the integer 2/6 discrete wavelet transform (DWT) filter pair described in [7] (and referred to as “filter A” in [1,8]).

II. ICER Overview

We start with a brief overview of some relevant concepts from the basic (2-D) ICER.

In ICER, multiple stages of a 2-D wavelet transform are applied to the image. The first stage is applied to the whole image, while subsequent stages are applied only to the (horizontally and vertically) low-pass subband from the previous stage. This results in the pyramidal decomposition first suggested by Mallat [9] and currently in common use. The resulting subbands include one small low-pass subband and several subbands that are high-pass in at least one dimension. A three-level 2-D wavelet decomposition of an image is shown in Fig. 1.

To limit the effect of data losses that can occur in transmission of data to Earth, ICER partitions image data into a user-defined number of error-containment segments, which are compressed independently. These segments are defined in the transform domain, and each segment approximately corresponds to

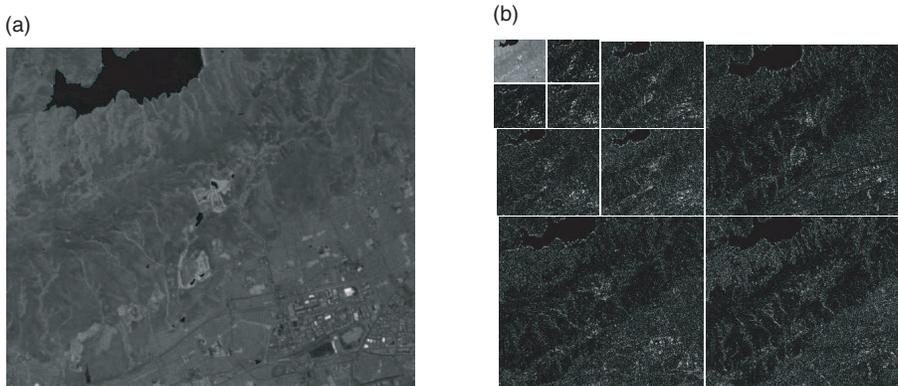


Fig. 1. Example of a three-level, 2-D wavelet decomposition: (a) is the original image, spectral band 41 (wavelength 740 nm) from the Moffett Field scene, and (b) is the result of the wavelet decomposition. In (b), in all subbands except the low-pass subband, absolute values are shown, contrast-enhanced by a factor of 3 relative to the low-pass subband.

³ The Moffett Field data set is available from the AVIRIS web site, <http://aviris.jpl.nasa.gov/html/aviris.freedata.html>.

a rectangular region of the original image. Figure 2 illustrates this correspondence. Note that the partitioning into segments is performed automatically based on the image dimensions and number of segments requested; this operation has no relation to the concept of segmentation for distinguishing objects or regions in an image. Segments are analogous to “precincts” in JPEG2000 [10].

Subbands that are high-pass in at least one dimension typically contain transform coefficients with a distribution that is roughly symmetric, has zero mean, and has a single sharp peak at zero. An approximately Laplacian distribution is prototypical.

In ICER, DWT coefficients are converted to sign-magnitude form and encoded one bit plane at a time starting with the most significant magnitude bit plane. (A bit plane is formed by taking the i th most significant magnitude bit of each coefficient of a subband, for some i .) When the first ‘1’ bit of a coefficient is encoded, the sign bit is encoded immediately afterward. If encoding is stopped after completing some number of bit planes, the resulting effective quantization is uniform except for a central deadzone. Bit-plane encoding is known to be an effective method for progressive compression of values for which the distribution has a sharp peak at zero; e.g., see [11]. A common non-progressive method is to quantize the coefficients and encode them in one pass. Either of these methods can be very effective, especially when combined with predictive coding and context modeling.

The low-pass subband resembles a low-resolution version of the original image and thus its DWT coefficient distribution can vary significantly from image to image. In ICER, for each error-containment segment of the low-pass subband a mean value is computed and subtracted (and encoded in the compressed bitstream). The resulting coefficient distribution has zero mean, but in general does not have a sharp peak near zero. The values are converted to sign-magnitude form and encoded one bit plane at a time as in the other subbands. The strong correlation between adjacent coefficients is exploited via predictive coding and context modeling. Other effective coding methods are possible.

III. Effect of Band-Dependent Signal Level Variations on 3-D Wavelet-Based Compression

In the straightforward extension of ICER to 3-D hyperspectral data sets, multiple stages of a 3-D wavelet transform are applied, with stages after the first applied only to the (spatially and spectrally) low-pass subband from the previous stage. The resulting decomposition is a 3-D Mallat decomposition (see Fig. 3); it is analogous to the 2-D Mallat decomposition of Fig. 1. Error-containment segments are

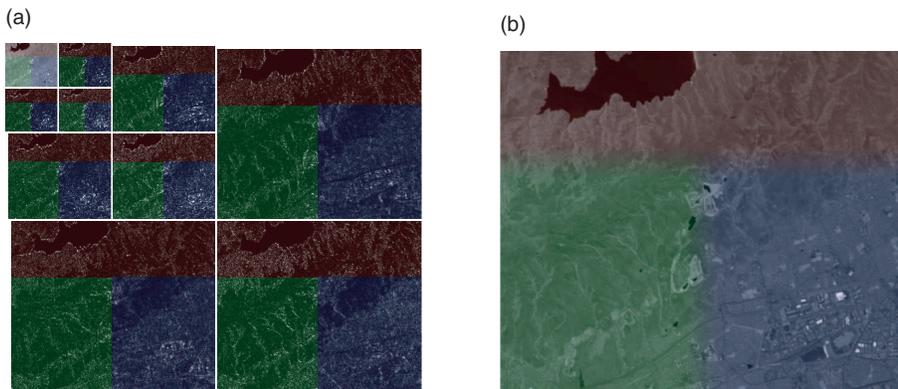


Fig. 2. The image of Fig. 1 divided into three error-containment segments, each tinted a different color: (a) shows the regions with hard boundaries in the transform domain, while (b) shows the resulting regions with soft boundaries in the original image.

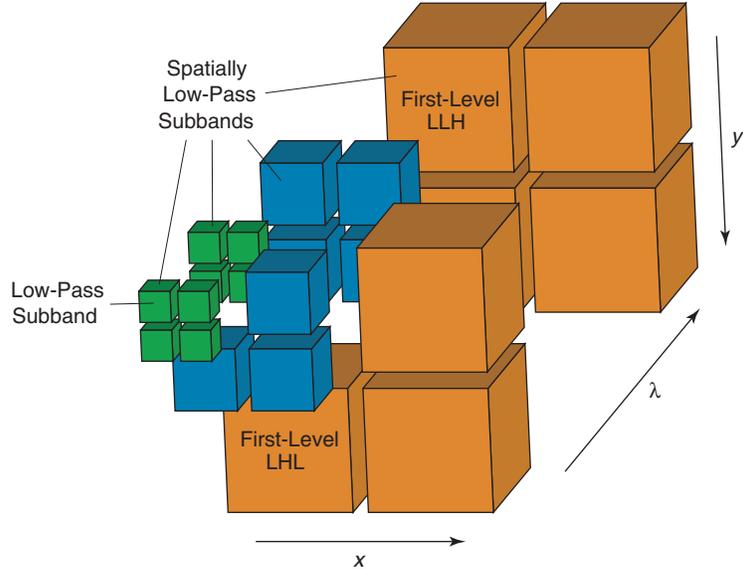


Fig. 3. The 3-D Mallat wavelet decomposition, illustrated here with three levels of decomposition. The x , y , and λ labels identify the horizontal, vertical, and spectral axes, respectively.

defined spatially (in the wavelet transform domain) so segments extend through all spectral bands. In our baseline implementation, in each segment of the low-pass subband the mean value is computed and subtracted.

In this section, for our examples we make use of the first level LLH (spatially low-pass, spectrally high-pass) subband and the first level LHL (horizontally low-pass, vertically high-pass, spectrally low-pass) subband.

A. Distributions of DWT Coefficients in Planes of Subbands

In any subband that is high-pass in at least one dimension, the mean value of the DWT coefficients will tend to be close to zero (see the Appendix). However, in a subband that is high-pass in only one dimension, individual planes that are orthogonal to the high-pass filter direction, such as xy planes (spatial planes) of the first level LLH subband and $x\lambda$ planes (horizontal-spectral planes) of the first level LHL subband, do not necessarily have mean values that are close to zero. In both of these subbands, the overall mean value is approximately zero, but in the LLH subband the individual xy planes have mean values that are far from zero, while in the LHL subband the mean values of the individual $x\lambda$ planes turn out to be much closer to zero. Figure 4 illustrates this situation for the Moffett Field scene. The overall histogram of the LLH subband as well as histograms of two individual spatial planes are shown in Fig. 5(a). Note from Fig. 5(b) that the overall histogram for the LHL subband is well-behaved. Comparing Fig. 5(a) and Fig. 6, we see that the LLH subband has a much narrower distribution (and consequently a higher peak) after subtracting the mean value from each spatial plane.

The widely varying mean values of spatial planes of the first level LLH subband are easily explained. The explanation applies to any subband that is spatially low-pass. The underlying cause is systematic differences in the signal level in different spectral bands. To a (very rough) first approximation, the spectra at individual spatial points are all similar (Fig. 7) due to effects such as atmospheric absorption in some regions of the spectrum and illumination from the same source (sunlight). Therefore, applying a wavelet decomposition in the spectral dimension results in similar transformed spectra (Fig. 8). In some spatial planes of the subband, the systematic content of the transformed spectra swamps the spatial

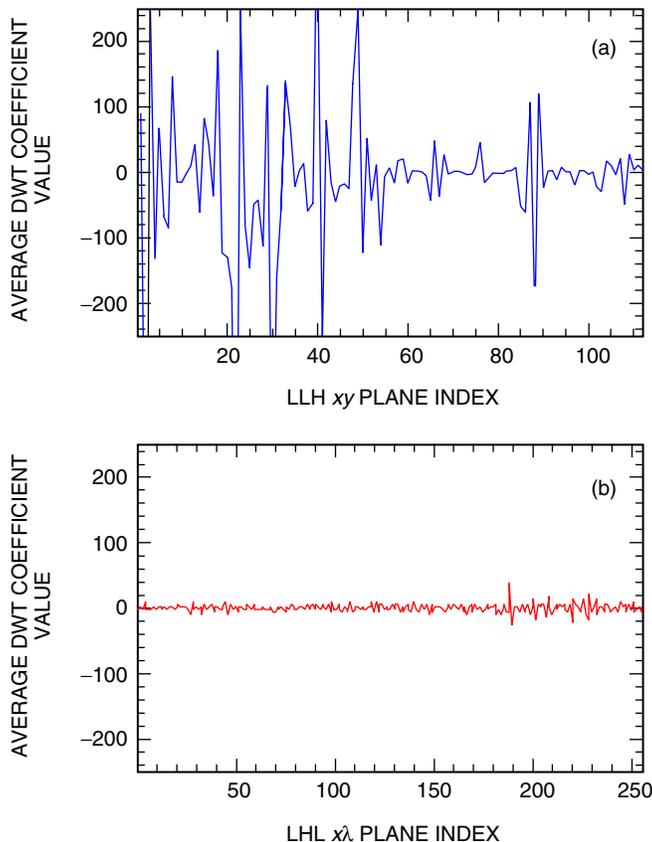


Fig. 4. Means of (a) the individual xy planes in the first level LLH subband and (b) the individual $x\lambda$ planes in the first level LHL subband. For both cases, the dataset is the Moffett Field scene.

variation that arises from the scene content; in other words, the magnitude of the mean value of some planes is relatively large compared to the variation of DWT coefficient values in the plane. High-pass filtering in either spatial direction effectively removes this systematic variation, so it is only an issue in spatially low-pass subbands.

The planes corresponding to spatial rows and columns of a hyperspectral image generally should not exhibit significant systematic differences in the signal level; thus, no analogous issues should arise with wavelet transforms in the spatial dimensions.

For a similar reason, an analogous effect generally does not arise in wavelet-based compression of 2-D images. However, it is instructive to consider the case of a 2-D image that has systematic variations in pixel intensities that depend on the row index or column index. Such variations do not occur in most types of images, but they are exhibited in a spatial-spectral plane of a hyperspectral image. An example of such a spatial-spectral plane is given in Fig. 9. Figure 10 shows the result of a 2-D wavelet decomposition of this spatial-spectral plane. Note that the systematic signal level differences in the spectral bands produce bright lines in the spatial (vertical) dimension in the spatially low-pass, spectrally high-pass subbands; these lines correspond to columns of these subbands that have mean values that are not close to zero.

We remark that, although in this case there are also significant lines in the spectral dimension of the spatially high-pass, spectrally low-pass subbands, the analogous phenomenon in the 3-D case is much weaker: in the 2-D case, these lines can be produced by a prominent spatial feature at a single location,

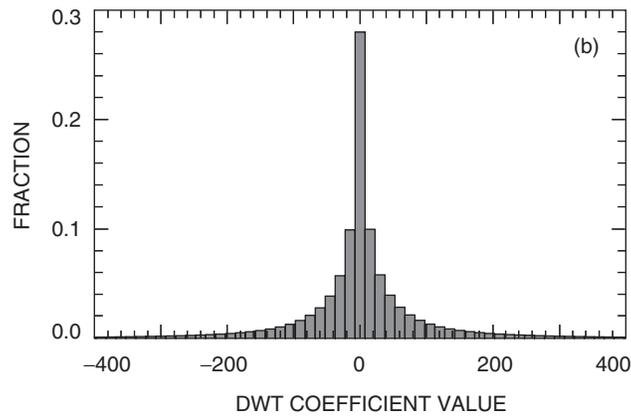
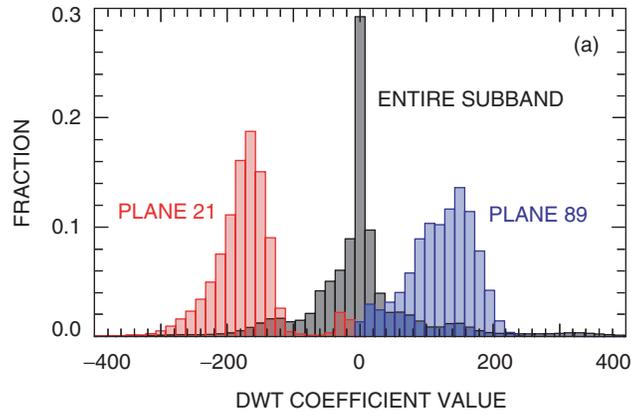


Fig. 5. Histograms of DWT coefficient values in subbands from the Moffett Field scene: (a) the first level LLH subband and two individual spatial planes of this subband and (b) the first level LHL subband.

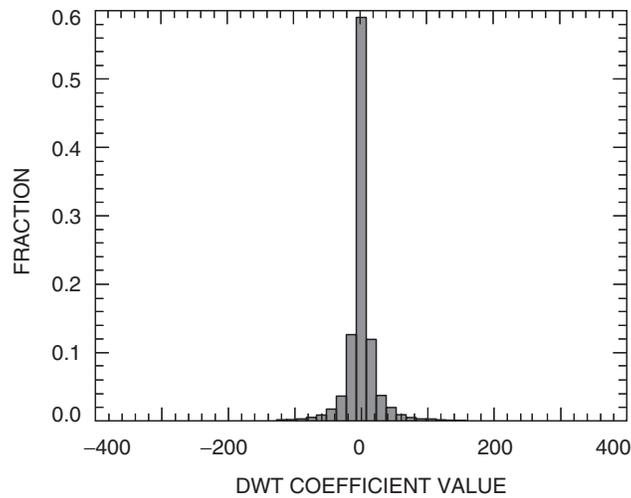


Fig. 6. Histogram of DWT coefficient values in the first level LLH subband from the Moffett Field scene after subtracting the mean value from each spatial plane.

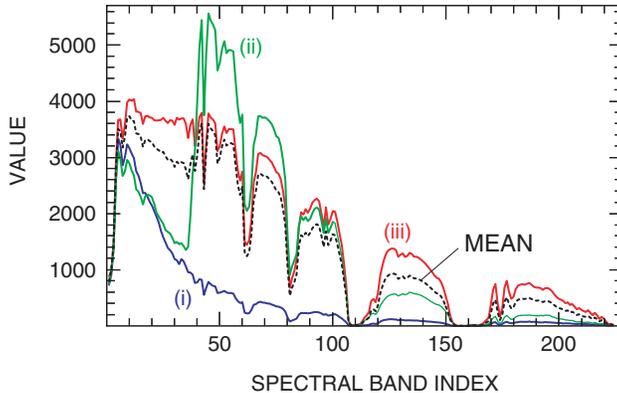


Fig. 7. Examples of individual spectra, labeled (i)–(iii), and the overall mean spectrum, from the Moffett Field scene.

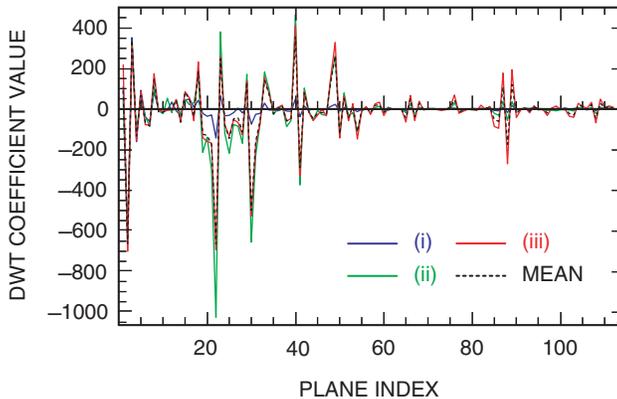


Fig. 8. The high-pass portion of a single 1-D wavelet decomposition of the individual spectra shown in Fig. 7, and of the mean spectrum for the whole scene.

but in the 3-D case it would take a prominent feature in a whole row of spatial locations to create a similar effect in an entire $x\lambda$ plane of an LHL subband. Thus, in a sense, there is more asymmetry among the dimensions in 3-D hyperspectral images than Fig. 10 suggests.

B. Challenges in Compressing Spatially Low-Pass Subbands

The bit-plane coding schemes used by ICER and other progressive compressors are best suited for DWT coefficient distributions that have mean zero and a single sharp peak at zero. It is implicitly assumed that the more significant magnitude bits are likely to be zero, and that after the most significant ‘1’ bit, the values of less significant bits are difficult to predict. These assumptions are appropriate for such distributions. Furthermore, sign-magnitude bit-plane coding schemes effectively produce quantization with a reconstruction point at zero, so even when few bit planes are encoded, the resulting coarse quantization can result in relatively low distortion since many values that are already very close to zero are quantized to zero.

We assert that, in the 3-D wavelet decomposition of a hyperspectral image, many spatial planes of spatially low-pass subbands have DWT coefficient distributions that are not well matched to ICER’s bit-plane coding scheme. As a result, the compression effectiveness suffers. In particular, for planes with

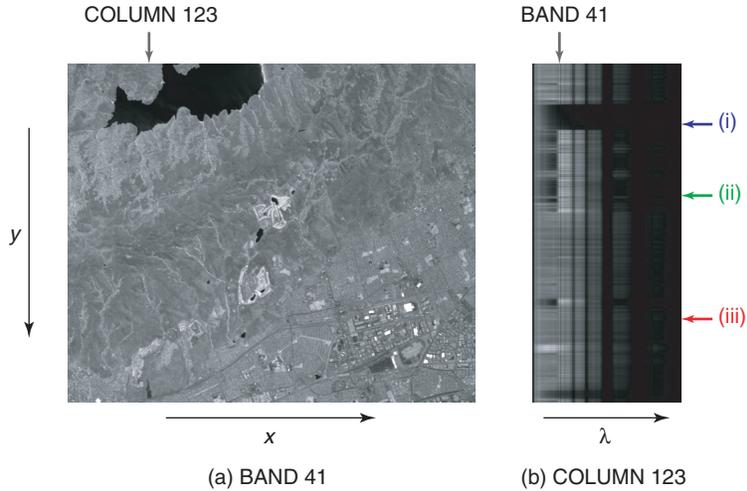


Fig. 9. Planes of the Moffett Field scene: (a) spectral band 41 and (b) column 123. Column 123 is the spatial–spectral plane consisting of column 123 from each spectral band. The arrows above the images indicate where they intersect. The labels (i)–(iii) identify the individual spectra used in Figs. 7 and 8.

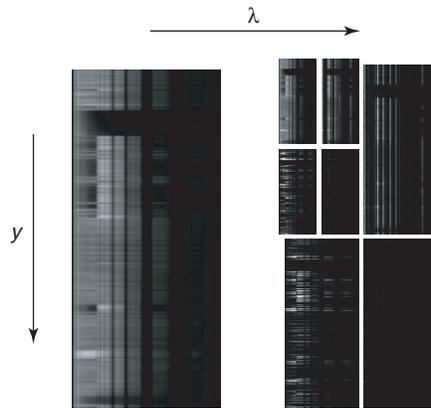


Fig. 10. The spatial–spectral plane of Fig. 9 and its two level 2-D wavelet decomposition. In all subbands except the low-pass subband, absolute values are shown, contrast-enhanced by a factor of 8 relative to the low-pass subband.

distributions that have a peak away from zero, the most significant ‘1’ bits occur earlier, and subsequent (refinement) bits are somewhat predictable because they often describe coefficient values near the peak. The bit-plane coding model does not adequately capture either of these effects. As a result, many bits are spent early for encoding spatially low-pass, spectrally high-pass subbands; this encoding is not as effective as it could be, hurting rate-distortion performance, especially at low bit rates. In addition, coarse quantization can result in relatively large distortion and a bias in the plane if there does not happen to be a reconstruction point near the peak value. Any bias in these planes can become a bias in several consecutive spectral bands in the reconstructed hyperspectral image. We give examples illustrating such biases in Section V. These effects would occur not only with ICER’s bit-plane coding scheme, but also with the schemes used by the Embedded Zerotree Wavelet (EZW) algorithm [12], Set Partitioning in Hierarchical Trees (SPIHT) [13], Embedded Block Coding with Optimized Truncation (EBCOT) [14], and JPEG2000 [10].

The biases described above can be considered to be a manifestation of a phenomenon we call *spectral ringing*. This term comes from the term “ringing,” which describes the more familiar 2-D image compression phenomenon in which an edge in an image produces spurious oscillations adjacent and parallel to the edges in the reconstructed image due to quantization effects. In hyperspectral data, the individual spectral bands play the role of edges and the resulting ringing occurs in the spectral dimension. We note that any spectral transform that does not take into account the systematic variations in the relative signal levels of the spectral bands would produce transform coefficient planes with systematic large-magnitude means, creating the possibility for analogous spectral compression artifacts.

Finally, we point out that the individual spatial planes of spatially low-pass subbands retain much of the qualitative appearance that is present in the original spectral bands. This contrasts with spatial planes of other subbands. Representative examples of spatial planes from the first level LHL and LLH subbands are shown in Fig. 11. The structure present in the spatial planes of spatially low-pass subbands suggests that exploiting the remaining correlation is important for effective compression.

IV. Handling Band-Dependent Signal Level Variations

In the preceding section, we saw that systematic variations in signal levels of different spectral bands can cause widely varying mean values in spatial planes of spatially low-pass subbands. We saw that this can have detrimental effects on image compression. We now describe two methods of mitigating these effects. Compression results illustrating the benefits of these methods are presented in Section V.

The two methods are as follows:

- (1) **Mean Subtraction.** The basic idea of this method is simply to subtract the mean values from spatial planes of spatially low-pass subbands prior to encoding, thus compensating for the fact that such spatial planes often have mean values that are far from zero. The resulting data (e.g., see Fig. 6) are better suited for compression by methods that are effective for subbands of 2-D images.
- (2) **Modified Decomposition.** Under this method, the subband decomposition is changed from the 3-D Mallat decomposition so that in stages of decomposition after the first, not only is the low-pass subband further decomposed, but spatially low-pass, spectrally high-pass subbands are also further decomposed spatially. An illustration of this subband decomposition is provided in Fig. 12; it should be compared to the Mallat decomposition shown in Fig. 3. The decomposition can be alternatively described as follows: first, a 2-D Mallat decomposition with the desired number of levels is performed (spatially) on every spectral band. Then, a single level of spectral decomposition is applied across the first level spatial subbands; a two-level 1-D Mallat decomposition is applied spectrally across the second level spatial subbands; and so on.

These two methods can be combined: we can perform the modified decomposition and then subtract the mean values from spatial planes of the spatially low-pass subbands.

In the context of ICER-3D, the mean subtraction method is implemented as follows. After the 3-D wavelet decomposition is performed, mean values are computed for and subtracted from each spatial plane of each error-containment segment of each spatially low-pass subband. The resulting data are converted to sign-magnitude form and compressed as in the baseline ICER-3D. The mean values are encoded in the compressed bitstream and added back to the data at the appropriate decompression step. The overhead incurred by encoding the mean values is only a few bits per spectral band per segment, which is negligible because of the huge size of hyperspectral data sets.

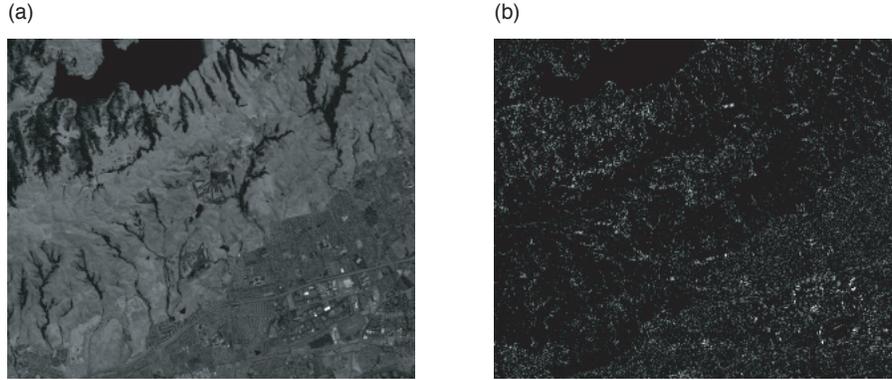


Fig. 11. Spatial planes of (a) the first level LLH and (b) the first level LHL subbands from the Moffett Field scene. Absolute values of DWT coefficients are shown. In both cases plane 89 (of 112) is shown.

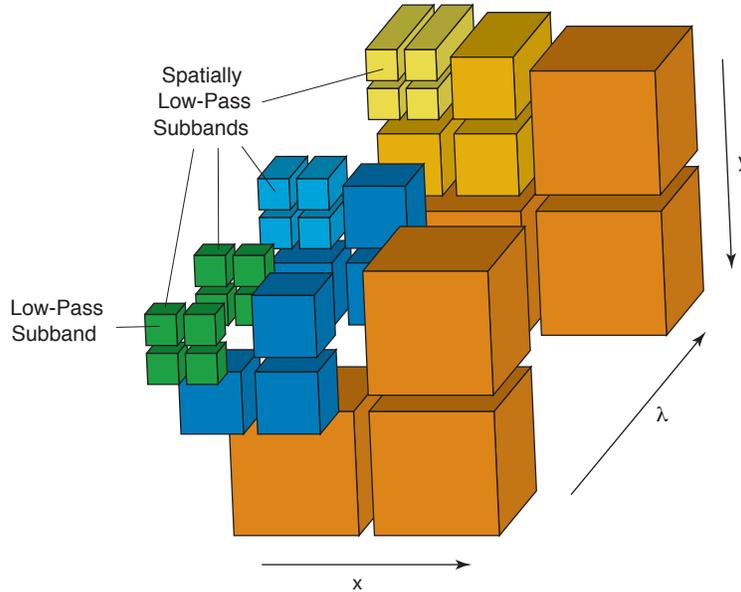


Fig. 12. The 3-D wavelet decomposition scheme used by the modified decomposition method, illustrated here with three levels of decomposition.

Note that it is important to subtract the means *after* all stages of subband decomposition; otherwise, if two adjacent error-containment segments have significantly different means, a sharp edge would appear after subtracting the means, artificially increasing high-frequency signal content in further stages of spatial decomposition.

The mean subtraction method is easy to implement. However, when used with the Mallat decomposition it has some tendency to produce visible segment boundaries in some reconstructed spectral bands at low bit rates (see Fig. 16 in Section V). When adjacent segments in a spatial plane of a spatially low-pass subband have different mean values, the subtraction of means causes the compressor to effectively use different quantizers, which can make the boundary between the segments conspicuous. Segment boundaries are generally not visible when mean values are subtracted only for the low-pass subband, as is the case for the baseline ICER-3D.

The modified decomposition is motivated by the observation that in a 3-D wavelet decomposition, the spatially low-pass, spectrally high-pass subbands have spatial planes that look qualitatively similar to spatial planes in the low-pass subband, as demonstrated in Fig. 11. This suggests that compression effectiveness improves with additional decompositions of the spatially low-pass, spectrally high-pass spatial planes, as it does with additional decompositions of the low-pass subband in a 2-D Mallat decomposition.

As we'll see in Section V, the mean subtraction and modified decomposition methods provide similar improvements in rate-distortion performance, but the latter appears to have a slight advantage in the subjective appearance of individual reconstructed bands. This is explained as follows. Because the spatially low-pass subbands are relatively small under the modified decomposition, they can be encoded to high fidelity with relatively few bits (and reasonable coding schemes will give these subbands high priority). As a result, low-pass data in the individual spectral bands will tend to be reproduced with high fidelity. Thus, the mean values of small regions of the reconstructed spectral bands will tend to be close to the corresponding values in the original image (where the size of the "small regions" increases with the number of levels of wavelet decomposition). The net result is that, even at fairly low bit rates, no noticeable bias will be present in small regions of the individual reconstructed spectral bands.

We note that, when the two methods are combined under ICER-3D, the fact that spatially low-pass subbands are represented with high fidelity means that segment boundaries will not be readily visible. Furthermore, whatever segment boundaries might be visible in the low-pass versions of the individual spectral bands will tend to be washed out by the blurring effect of several levels of the inverse wavelet transform.

Other researchers have also devised hyperspectral image compression schemes that use 3-D wavelet decompositions that are modifications to the Mallat decomposition. For example, in [3,4] the wavelet decomposition used is equivalent to a 2-D Mallat decomposition in the spatial domain followed by a 1-D Mallat decomposition in the spectral dimension. The resulting overall decomposition has further decomposed subbands compared to our modified decomposition with the same number of stages. Because all of the transform steps of our modified decomposition are included in the decomposition of [3,4], the latter enables a similar advantage in compression effectiveness. Alternatives to the Mallat 3-D wavelet decompositions have also been used for compression of 3-D medical data sets (e.g., [15]) and video coding (e.g., [16,17]). In video coding, the number of samples in the temporal dimension has generally been very small.

Finally, we outline an approach to encoding spatially low-pass subbands that is an alternative to the mean subtraction method. As we have observed above, spatial planes of spatially low-pass subbands look rather image-like, and the distributions of DWT coefficients in such spatial planes often do not have zero means and single sharp peaks. Rather than adjusting the effective quantizers by subtracting mean values, one could alter the way these subbands are encoded so as to be based on (possibly progressive) uniform quantization that does not depend on subband content. One method of doing this is bit-plane coding of the subbands without mean subtraction or conversion to sign-magnitude form. For effective compression under this method, one would want to use predictive coding and context modeling modules that can effectively exploit the image-like appearance of the spatial planes of these subbands. This modification could be applied to either the standard Mallat decomposition or the modified decomposition.

Since all error-containment segments in a subband would be using the same quantizer under this alternative encoding approach (so long as they are all compressed to the same fidelity), it eliminates any possible boundary artifacts between segments. However, this coding approach might not encode the affected (i.e., spatially low-pass) subbands as effectively. In particular, under the mean subtraction method a sharp peak in the DWT coefficient distribution would tend to be exploited by quantization that includes a reconstruction point at the peak value, but under uniform quantization, having a reconstruction point near the peak values would be less likely, and as a result there could be higher distortion and a systematic bias in some spectral bands of the reconstructed image.

V. Results

The methods described in Section IV provide a noticeable improvement in rate-distortion performance compared to the baseline approach, especially at moderate to low bit rates (roughly 1 bit/pixel/band and below). In Fig. 13, we compare the rate-distortion performance of these methods to the baseline approach for the Moffett Field scene and for a 512-line radiance data scene of Arizaro, Argentina, taken on February 7, 2001.⁴ The points shown on the curves were produced by encoding all subband bit planes that have significance exceeding a given value. It is seen that mean subtraction, modified decomposition, and modified decomposition with mean subtraction

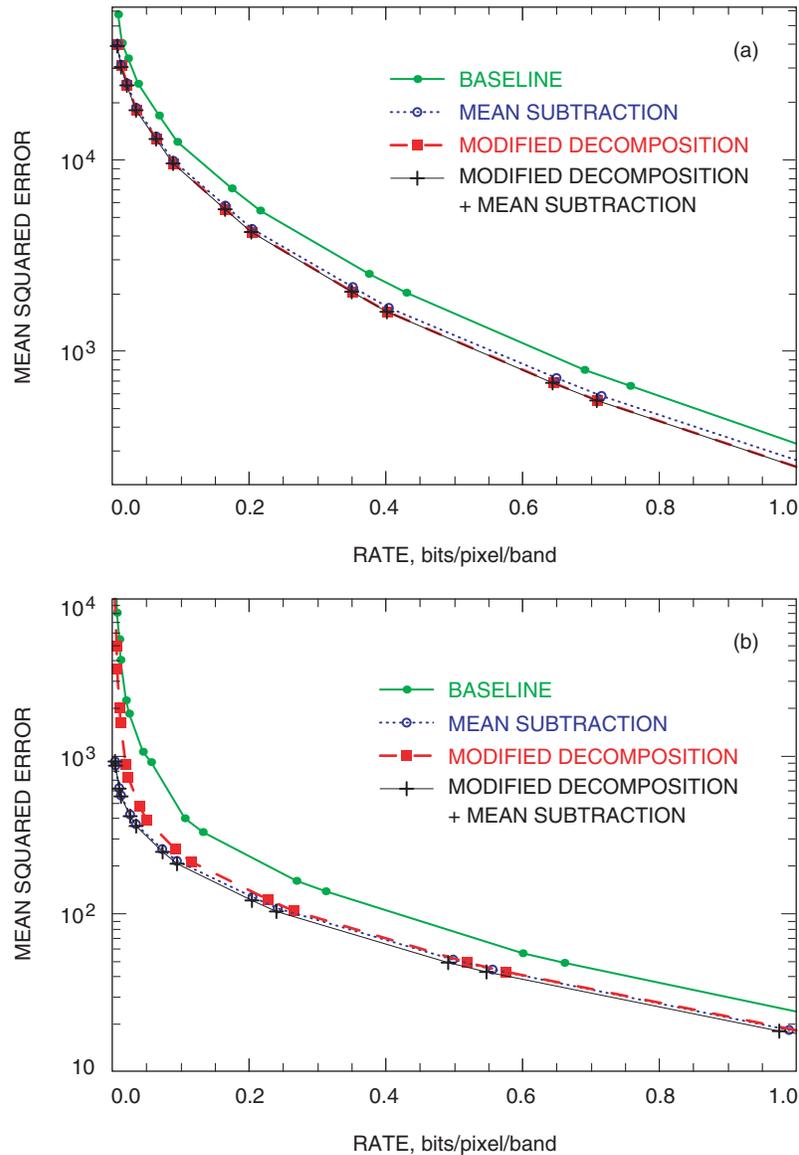


Fig. 13. Rate-distortion performance of the methods described in Section IV and baseline ICER-3D for (a) the Moffett Field scene using 5 stages of wavelet decomposition and a single error-containment segment and (b) the Arizaro scene using 3 stages of wavelet decomposition and 4 error-containment segments.

⁴ AVIRIS flight number f010207t01p02_r06.

and the combination of the two provide very similar rate-distortion performance, and, for example, give roughly a 10 percent improvement in rate compared to the baseline method at 1 bit/pixel/band. When the number of wavelet decompositions is small, the rate-distortion performance of the modified decomposition alone is slightly worse than that of mean subtraction alone or the combination of the two methods.

Overall, the use of either method from Section IV with ICER-3D provides a moderate subjective image quality improvement consistent with the improvement in mean squared error (MSE) distortion. In some cases, however, the improvement is more dramatic, especially with regard to reduction of bias in reconstructed spectral bands when compressed at low bit rates. This is illustrated in the false-color images of Figs. 14 through 16. These images were produced by mapping band 176 to red, 81 to green, and 33 to blue. Bands 176 and 81 were chosen because their reconstructions exhibit a noticeable bias when using the baseline ICER-3D. This bias appears as an apparent overall color change in the false-color images. To a lesser degree, regional biases can be seen under the mean subtraction method.

As the discussion in Section IV suggests, error-containment segment boundaries are sometimes visible when mean subtraction is used alone. This is illustrated in the reconstruction of Fig. 16, where some segment boundaries near the top of the image are noticeable under mean subtraction but not under the other methods.

In Figs. 14 through 16, reconstructions using the modified decomposition combined with mean subtraction are not shown because they are visually indistinguishable from those produced by the modified decomposition without mean subtraction.

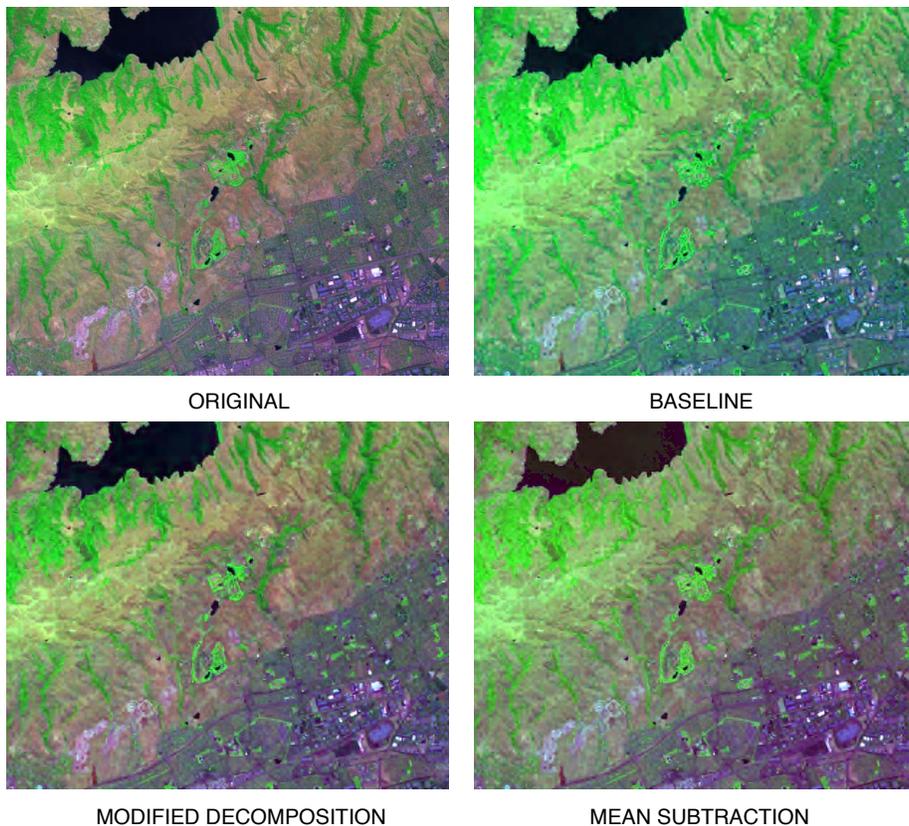


Fig. 14. Reconstructed false-color images of the Moffett Field scene using the baseline ICER-3D, and ICER-3D with mean subtraction and modified decomposition. In all three cases, the entire hyperspectral scene was compressed to 0.1 bits/pixel/band using 5 stages of wavelet decomposition and a single error-containment segment.

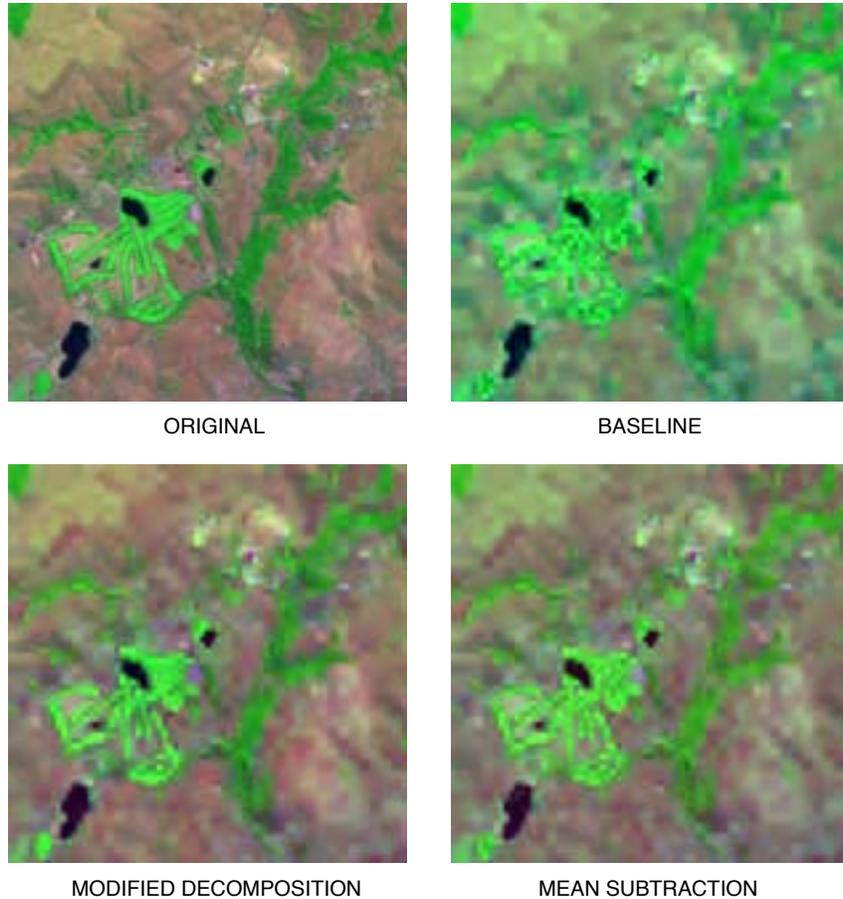


Fig. 15. Detail region from the reconstructed false-color images of Fig. 14.

VI. Conclusion

When using 3-D wavelet transforms for hyperspectral image compression, systematic variations in signal level of different spectral bands can cause widely varying mean values in spatial planes of spatially low-pass subbands. Failing to account for this phenomenon can have detrimental effects on image compression, including reduced effectiveness in compressing spatially low-pass subband data, and biases in some reconstructed spectral bands.

These effects can be mitigated by subtracting the mean value from each spatial plane of each spatially low-pass subband, or by modifying the wavelet decomposition to perform extra stages of spatial decomposition in spatially low-pass subbands, or by a combination of these approaches. We presented examples illustrating that these methods offer similar improvements in rate-distortion performance. Both approaches reduce biases in reconstructed spectral bands and provide an improvement in subjective reconstructed image quality. The modified decomposition has the advantage that it does not have a tendency to produce visible boundaries between error-containment segments.

Acknowledgment

The authors would like to thank Robert Green for providing the Arizaro image used in Section V.

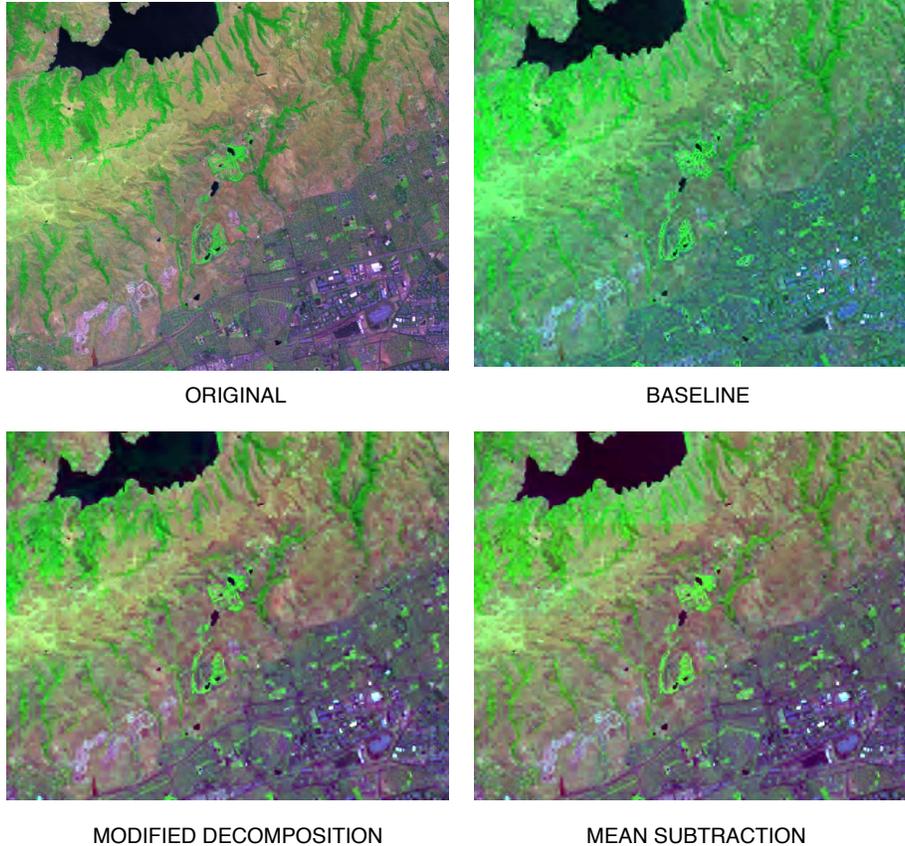


Fig. 16. Reconstructed false-color images of the Moffett Field scene using the baseline ICER-3D, and ICER-3D with mean subtraction and modified decomposition. In all three cases, the entire hyperspectral scene was compressed to 0.0625 bits/pixel/band using 5 stages of wavelet decomposition and 16 error-containment segments.

References

- [1] A. Kiely and M. Klimesh, "The ICER Progressive Wavelet Image Compressor," *The Interplanetary Network Progress Report 42-155, July–September 2003*, Jet Propulsion Laboratory, Pasadena, California, pp. 1–46, November 15, 2003. http://ipnpr.jpl.nasa.gov/tmo/progress_report/42-155/155J.pdf
- [2] A. Kiely and M. Klimesh, "Preliminary Image Compression Results from the Mars Exploration Rovers," *The Interplanetary Network Progress Report*, vol. 42-156, Jet Propulsion Laboratory, Pasadena, California, pp. 1–8, February 15, 2004. http://ipnpr.jpl.nasa.gov/tmo/progress_report/42-156/156I.pdf
- [3] X. Tang, S. Cho, and W. A. Pearlman, "3D Set Partitioning Coding Methods in Hyperspectral Image Compression," *Proc. 2003 International Conference on Image Processing*, vol. II, pp. II-239–II-242, September 14–17, 2003.
- [4] Y. Wang, J. T. Rucker, and J. E. Fowler, "Three-Dimensional Tarp Coding for the Compression of Hyperspectral Images," *IEEE Geoscience and Remote Sensing Letters*, vol. 1, no. 2, pp. 136–140, April 2004.

- [5] S. Lim, K. Sohn, and C. Lee, "Compression for Hyperspectral Images Using Three Dimensional Wavelet Transform," *Proc. IEEE 2001 International Geoscience and Remote Sensing Symposium (IGARSS '01)*, vol. 1, pp. 109–111, 2001.
- [6] G. Vane, R. Green, T. Chrien, H. Enmark, E. Hansen, and W. Porter, "The Airborne Visible/Infrared Imaging Spectrometer (AVIRIS)," *Remote Sensing of Environment*, vol. 44, pp. 127–143, 1993.
- [7] M. D. Adams and F. Kossentini, "Reversible Integer-to-Integer Wavelet Transforms for Image Compression: Performance Evaluation and Analysis," *IEEE Transactions on Image Processing*, vol. 9, no. 7, pp. 1010–1024, June 2000.
- [8] A. Said and W. Pearlman, "An Image Multiresolution Representation for Lossless and Lossy Compression," *IEEE Transactions on Image Processing*, vol. 9, no. 5, pp. 1303–1310, September 1996.
- [9] S. G. Mallat, "A Theory for Multiresolution Signal Decomposition: The Wavelet Representation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 11, no. 7, pp. 674–693, July 1989.
- [10] D. S. Taubman and M. W. Marcellin, *JPEG2000: Image Compression Fundamentals, Standards and Practice*, Boston, Massachusetts: Kluwer Academic Publishers, 2002.
- [11] A. Kiely, "Progressive Transmission and Compression of Images," *The Telecommunications and Data Acquisition Progress Report 42-124, October–December 1995*, Jet Propulsion Laboratory, Pasadena, California, pp. 88–103, February 15, 1995. http://tmo.jpl.nasa.gov/tmo/progress_report/42-124/124E.pdf
- [12] J. M. Shapiro, "Embedded Image Coding Using Zerotrees of Wavelet Coefficients," *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3445–3462, December 1993.
- [13] A. Said and W. Pearlman, "A New, Fast, and Efficient Image Codec Based on Set Partitioning in Hierarchical Trees," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 6, no. 3, pp. 243–250, June 1993.
- [14] D. Taubman, "High Performance Scalable Image Compression with EBCOT," *IEEE Transactions on Image Processing*, vol. 9, no. 7, pp. 1158–1170, July 2000.
- [15] Z. Xiong, X. Wu, S. Cheng, and J. Hua, "Lossy-to-Lossless Compression of Medical Volumetric Data Using Three-Dimensional Integer Wavelet Transforms," *IEEE Transactions on Medical Imaging*, vol. 22, no. 3, pp. 459–470, March 2003.
- [16] B.-J. Kim, Z. Xiong, and W. A. Pearlman, "Low Bit-Rate Scalable Video Coding with 3-D Set Partitioning in Hierarchical Trees (3-D SPIHT)," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 10, no. 8, pp. 1374–1387, December 2000.
- [17] J. Y. Tham, S. Ranganath, and A. A. Kassim, "Highly Scalable Wavelet-Based Video Codec for Very Low Bit-Rate Environment," *IEEE Journal on Selected Areas in Communications*, vol. 16, no. 1, pp. 12–27, January 1998.

Appendix

Mean Value of High-Pass DWT Coefficients

In this Appendix, we provide some justification for the usual assumption that a subband that is high-pass in at least one direction will have a mean DWT coefficient value that is close to zero. We do this by deriving an approximate expression for the mean high-pass coefficient value of a single decomposition in the one-dimensional case.

Consider a length N signal x_1, \dots, x_N . For convenience we assume N is even. For a fairly general linear DWT, a high-pass coefficient h_n can be computed as $h_n = \sum_i c_i x_{2n+i}$, where the sum is over the indices of the wavelet filter coefficients c_i . Here we have assumed that the number of filter coefficients is small, and n is in the range $1, \dots, N/2$ but not so near the edge of that range that boundary effects matter. Because these are high-pass values, the sum of the filter coefficients c_i should be zero.

To get a reasonably simple approximate expression for the mean high-pass coefficient value \bar{h} , we make the approximation that each x_j makes the same contribution to \bar{h} . Specifically,

$$\bar{h} = \frac{2}{N} \sum_{n=1}^{N/2} h_n \approx \frac{2}{N} \sum_n \sum_i c_i x_{2n+i}$$

where the last pair of sums is over those n and i for which $2n+i$ is in the range $1, \dots, N$. Separating the odd and even indices of x gives

$$\begin{aligned} \bar{h} &\approx \frac{2}{N} \left(\sum_{j \text{ odd}} x_j \sum_{i \text{ odd}} c_i + \sum_{j \text{ even}} x_j \sum_{i \text{ even}} c_i \right) \\ &= \bar{x}_{\text{odd}} \sum_{i \text{ odd}} c_i + \bar{x}_{\text{even}} \sum_{i \text{ even}} c_i \\ &= \bar{x}_{\text{odd}} \sum_i c_i + (\bar{x}_{\text{even}} - \bar{x}_{\text{odd}}) \sum_{i \text{ even}} c_i \\ &= (\bar{x}_{\text{even}} - \bar{x}_{\text{odd}}) \sum_{i \text{ even}} c_i \end{aligned}$$

where \bar{x}_{even} and \bar{x}_{odd} denote the mean even- and odd-indexed signal values, respectively, and we have used the fact that the c_i sum to zero.

In most situations, we would expect \bar{x}_{even} and \bar{x}_{odd} to be about equal due to the effect of averaging many samples, which implies that \bar{h} is close to zero. This in turn implies that a subband that is high-pass in at least one direction will have a mean DWT coefficient value that is close to zero.